

Claims

- [c1] 1. A method for extended memory support in a database system having a primary cache, the method comprising:
creating a secondary cache in memory available to the database system;
mapping a virtual address range to at least a portion of the secondary cache;
when the primary cache is full, replacing pages from the primary cache using the secondary cache;
in response to a request for a particular page, searching for the particular page in the secondary cache if the particular page is not found in the primary cache;
if the particular page is found in the secondary cache, determining a virtual address in the secondary cache where the particular page resides based on the mapping;
and
swapping the particular page found in the secondary cache with a page in the primary cache, so as to replace a page in the primary cache with the particular page from the secondary cache.
- [c2] 2. The method of claim 1, wherein said creating step includes creating the secondary cache using a shared

memory file system.

- [c3] 3. The method of claim 2, wherein the shared memory file system is available as part of an operating system on a computer platform on which the database system is running.
- [c4] 4. The method of claim 3, wherein the operating system comprises a Linux operating system.
- [c5] 5. The method of claim 1, wherein said mapping step includes using a memory mapped file function.
- [c6] 6. The method of claim 5, wherein the memory mapped file function is available as part of an operating system on a computer platform on which the database system is running.
- [c7] 7. The method of claim 1, wherein said creating step includes creating the secondary cache on external memory available to the database system.
- [c8] 8. The method of claim 1, wherein said swapping step includes consulting a least recently used (LRU) list maintained for the primary cache to determine the page to be moved to the secondary cache.
- [c9] 9. The method of claim 8, wherein said swapping step further comprises copying the page to be moved to the

secondary cache to a temporary buffer.

- [c10] 10. The method of claim 9, wherein said swapping step further comprises moving the particular page from the secondary cache to address of the page in the primary cache to be moved to the secondary cache.
- [c11] 11. The method of claim 9, wherein said swapping step further comprises moving the page from the temporary buffer to the secondary cache.
- [c12] 12. The method of claim 11, wherein further comprising: adding the replaced page to a most recently used end of a most recently used/least recently used (MRU/LRU) list maintained for the secondary cache.
- [c13] 13. The method of claim 1, wherein said replacing step includes maintaining a least recently used (LRU) list for the primary cache and selecting the page to be moved to the secondary cache based on said LRU list.
- [c14] 14. The method of claim 1, further comprising: providing a washing mechanism in the secondary cache for writing pages in the secondary cache to disk.
- [c15] 15. The method of claim 14, wherein a page is written from the secondary cache to disk in response to copying a page from disk to the primary cache.

- [c16] 16. The method of claim 15, wherein the page written from the secondary cache to disk is selected, based at least in part, on a most recently used/least recently used (MRU/LRU) list maintained for the secondary cache.
- [c17] 17. The method of claim 1, wherein said replacing step includes determining pages to be maintained in the secondary cache.
- [c18] 18. The method of claim 17, wherein said determining step includes determining pages to be maintained in the secondary cache based, at least in part, on workload of the database system.
- [c19] 19. The method of claim 1, wherein said replacing step includes substeps of:
moving a page from the primary cache to the secondary cache; and
reading a page into the primary cache from disk.
- [c20] 20. The method of claim 19, wherein said substep of moving a page from the primary cache includes selecting a page from the primary cache based on a least recently used (LRU) list maintained for the primary cache.
- [c21] 21. The method of claim 19, wherein said substep of moving a page from the primary cache includes selecting

a location for the page in the secondary cache based on a most recently used/least recently used (MRU/LRU) list maintained for the secondary cache.

- [c22] 22. A computer-readable medium having processor-executable instructions for performing the method of claim 1.
- [c23] 23. A downloadable set of processor-executable instructions for performing the method of claim 1.
- [c24] 24. A database system providing extended memory support, the system comprising:
a primary cache for maintaining data pages used by the database system in addressable memory available to the database system;
a secondary cache for maintaining data pages replaced from the primary cache in extended memory available to the database system;
a search module for receiving a request from a user for a particular data page and determining whether the particular data page is in secondary cache if the particular data page is not in the primary cache; and
a module for replacing a data page in the primary cache with the particular data page from the secondary cache if the particular data page is found in the secondary cache.

- [c25] 25. The system of claim 24, wherein the secondary cache is implemented using a shared memory file system.
- [c26] 26. The system of claim 25, wherein the shared memory file system is available as part of an operating system on a computer platform on which the database system is running.
- [c27] 27. The system of claim 26, wherein the operating system comprises a Linux operating system.
- [c28] 28. The system of claim 24, wherein the secondary cache is mapped to the extended memory using a memory mapped file function.
- [c29] 29. The system of claim 28, wherein the memory mapped file function is available as part of an operating system on a computer platform on which the database system is running.
- [c30] 30. The system of claim 24, wherein the secondary cache is created on external memory available to the database system.
- [c31] 31. The system of claim 24, wherein the primary cache includes a least recently used (LRU) list for determining data pages to be moved to the secondary cache.
- [c32] 32. The system of claim 31, wherein the module for re-

placing consults the LRU list for selecting the data page to be moved to the secondary cache.

[c33] 33. The system of claim 24, wherein the module for replacing copies the data page to be moved to the secondary cache to a temporary buffer.

[c34] 34. The system of claim 33, wherein the module for replacing moves the particular data page from the secondary cache to address of the data page in the primary cache to be moved to the secondary cache.

[c35] 35. The system of claim 33, wherein the module for replacing moves the data page from the temporary buffer to the secondary cache.

[c36] 36. The system of claim 35, wherein the secondary cache includes a most recently used/least recently used (MRU/LRU) list and the module for replacing adds the data page moved to the secondary cache to the most recently used end of said MRU/LRU list.

[c37] 37. The system of claim 24, further comprising:
a washing mechanism in the secondary cache for writing data pages in the secondary cache to disk.

[c38] 38. The system of claim 37, wherein the washing mechanism writes a data page in the secondary cache to disk in

response to copying a data page from disk to the primary cache.

[c39] 39. The system of claim 38, wherein the washing mechanism selects the data page from the secondary cache based, at least in part, on a most recently used/least recently used (MRU/LRU) list maintained for the secondary cache.

[c40] 40. The system of claim 24, wherein the module for replacing determines data pages to be maintained in the secondary cache.

[c41] 41. The system of claim 24, further comprising:
a module for reading a page into the primary cache from disk.

[c42] 42. The system of claim 41, wherein the module for reading selects a data page from the primary cache to be moved to the secondary cache if the primary cache is full.

[c43] 43. The system of claim 42, wherein the module for reading selects the data page based on a least recently used (LRU) list maintained for the primary cache.

[c44] 44. The system of claim 42, wherein the module for reading selects a location in the secondary cache for the

data page to be moved from the primary cache.

- [c45] 45. The system of claim 44, wherein the module for reading selects the location in the secondary cache based on a most recently used/least recently used (MRU/LRU) list maintained for the secondary cache.